

ECHO Client Partner User Guide - Chapter 1

- Before you begin
 - Tasks That You Will Perform as a Client Partner
 - Skills You Will Need as a Client Partner
 - ECHO Concept and Design
 - ECHO as a Spatially Enabled Metadata Search and Order System
 - Security
 - Supported Platforms
 - ECHO Capability and Functionality
 - ECHO Benefits to Client Partners
 - Ease of Participation
 - Cost to Field
 - Cost to Operate
 - Extensibility
- ECHO Systems

Guide Navigation

1. Preface
2. **Before you begin**
3. The basics
4. Logging in, setting up and getting started
5. Querying for Earth Science metadata
6. The ECHO Alternative Query Language
7. Ordering data through ECHO
8. Metadata subscriptions
9. Appendices

Before you begin

The primary reason for designing the EOS ClearingHouse (ECHO) was to increase access to Earth science data and services by providing a system with a machine-to-machine interface, that is, an Application Programming Interface (API).

ECHO functions as a metadata clearinghouse of Earth science metadata for a wide variety of partners, enabling the science community to exchange information. Data Partners provide the ECHO community with metadata representing their Earth science data holdings. ECHO technology in turn provides services for Client Partners and Data Partners and supports efficient discovery and access to Earth science data. ECHO also functions as an order broker for the data, and offers services applied to that data. ECHO provides a portal on the internet where ECHO clients can search the metadata for information they wish to order.

Client applications can access data holdings via order distribution or online access. Data Partners retain complete control over what metadata are represented in ECHO including inserting new metadata, modifying existing metadata and removing old metadata, and controlling access to their metadata.

Tasks That You Will Perform as a Client Partner

Usually performed in the order shown below:

In [Chapter 4](#)

- Logging in and getting started
 - Creating and managing ECHO sessions
 - Creating and managing user accounts

In [Chapters 5, 6 and 8](#)

- Querying for Earth Science Data
 - Formatting query results
 - Handling large result sets
 - Using subscriptions to automate queries

- Creating and deleting subscriptions
- Listing subscriptions

In [Chapter 7](#)

- Ordering data through ECHO
 - Adding order items
 - Updating order items
 - Removing order items
 - Removing orders
 - Setting user information for an order
 - Validating orders
 - Requesting a quote for an order
 - Submitting an order
 - Tracking or canceling an order

In [Chapter 8](#)

- Using Event Notification Service
 - Setting delivery modes
 - Filtering events
 - Creating event subscriptions
 - Removing event subscriptions
 - Renewing event subscriptions

Skills You Will Need as a Client Partner

Since ECHO uses platform-independent web service definitions for its API, there are no requirements for a client programming language. All examples in this document are in snippets of Java code; however, the code samples provided could be translated to any web service capable language.

As an ECHO Data Partner, you need to be familiar with basic software development and Service Oriented Architecture (SOA) concepts such as:

- XML and XML Schema (XSD)
- Client/Server-based programming (client stubs, remote endpoints, etc.)

As a REST-API user,

- A basic understanding of the REST concept
- HTTP

As a SOAP-API user,

- Web Service Definition Language (WSDL)
- Service-based Application Programmer's Interface (API)

ECHO Concept and Design

NASA's Earth Science Data and Information System (ESDIS) has built ECHO based on Extensible Markup Language (XML) and Web Service technologies. ECHO interfaces with different clients and users through its series of Application Program Interfaces (APIs). ECHO is an open system with published APIs available to the ECHO Development and User community.

ECHO is a middleware application, and interacting with ECHO means interacting with the ECHO API. There is typically a user-focused client application interacting with ECHO's API on behalf of an end user. This client may be a generic, query and order-based client, or may be specific to an end user's research, mission, or general area of interest. ECHO incorporates a Universal Description, Discovery, and Integration (UDDI) registry to facilitate registration, discovery, and invocation of services related to the ECHO holdings. Internally, ECHO specifies APIs and provides middleware components, including data and service search and access functions, in a layered architecture. The figure below depicts the ECHO system context in relation to its public APIs.

All ECHO metadata is stored in an Oracle database with spatial extensions. The metadata model is derived primarily from that used by the Earth Observing System Data and Information System (EOSDIS) Core System (ECS). For more details about the ECHO model, refer to [ECHO Earth Science Metadata Model chapter of the Data Partner Guide](#)

Key features of the ECHO architecture are:

- **Ease of Partner Participation** Designed to be low-cost and minimally intrusive, ECHO offers a set of standard ways for partners to interface with the system and a metadata exchange approach that accommodates existing partners and technology.

- **Data Model Consistency** To mitigate the risk of being unable to match all possible partner data models, ECHO has prototyped a Metadata Mapping Tool to translate non-standard formats upon ingest into ECHO.
- **Open System/Published APIs** To accommodate independent ECHO clients, ECHO uses an open system approach and publishes domain APIs. These APIs are independent of the underlying transport protocols used. ECHO communicates using WS-I Basic Profile v1.0 compliant web services. This API is located at api.echo.nasa.gov/echo. Interactions with ECHO may involve user interactions in real time or may be machine to machine.
- **Evolutionary Development** The ECHO system is being developed incrementally to allow for insight and feedback during the development cycle. Industry trends are followed and the use of commercial, off-the-shelf (COTS) products is optimized.
- **Extensibility of Client User Interfaces and Capabilities** ECHO extensibility is ensured by its component architecture, which allows new capabilities and functions to be plugged in, modeling relationships between services/APIs/UIs, and continued prototyping. ECHO's current focus is on the middleware and on enabling many different types of user interfaces via its APIs.

ECHO as a Spatially Enabled Metadata Search and Order System

Oracle enables the ECHO system to interact with spatially enabled Earth science metadata by use of spatial extensions into the system and business logic within the system that understands how to interact with that metadata. In addition, a second ECHO interface (Ingest) allows metadata updates to go directly into the database, bypassing the message-passing API. The File Transfer Protocol (FTP) server is configured to receive these update files, which are expressed in XML conforming to three schemas, one for granules (or inventory), one for collections (or datasets), and one for browse.

Note: For ECHO 10.0 and later, these formats are schemas; for legacy ECHO, these formats are DTDs.

The schemas are defined on the ECHO 10.10 Ingest DTD/Schemas page of the ECHO website: <http://api.echo.nasa.gov/echo/apis.html>

Oracle's spatial capabilities support queries for ECHO metadata whose spatial extent is described within the system. A Data Partner can define the spatial extent of a granule or a collection with different spatial constructs (for example: point and polygon). A Client Partner can then construct a search using a point, a line, or a polygon (or multiple polygon) spatial type, and ECHO responds with data whose spatial region intersects the described region.

ECHO provides services for interacting with its **Catalog** of metadata. Queries can be performed in a number of ways; result formats can be specified, and the resulting data sets can be incrementally accessed so that large return sets can be handled gracefully. ECHO also supports constructing, submitting, and tracking orders for the data that the metadata represents. ECHO supports both an embedding of a Uniform Resource Locator (URL) within the metadata for accessing the data (which the client simply accesses via Hypertext Transfer Protocol [HTTP]), and a more complicated order process in which quotes and order options are accommodated.

ECHO incorporates the ECS concept of granules and collections and defines separate DTDs for updating each, under the assumption that granules will indicate which collection is considered their primary collection. Primary collection means the collection that owns the granule.

A **collection** is a grouping of granules that all come from the same source, such as a modeling group or institution. Collections have information that is common across all the granules they "own" and a template for describing additional attributes not already part of the metadata model.

A **granule** is the smallest aggregation of data that can be independently managed (described, inventoried, and retrieved). Granules have their own metadata model and support values associated with the additional attributes defined by the owning collection.

A third type of metadata, which is not spatially enabled but useful, is **browse metadata**, which provide a high-level view of granule or collection metadata and cross-referencing to other granules or collections.

Security

The ECHO system supports Secure Sockets Layer (SSL)-based communication, which a client must use to pass passwords or other sensitive information securely. Internally, the systems are firewalled to prevent unintended access.

Supported Platforms

The ECHO system supports clients capable of initiating an HTTP connection from a variety of programming languages

ECHO Capability and Functionality

ECHO provides an infrastructure that allows various communities to share tools, services, and metadata. As a metadata clearinghouse, it supports many data access paradigms such as navigation and discovery. As an order broker, ECHO forwards orders for data discovered through the metadata query process to the appropriate Data Partners for order fulfillment. As a service broker, ECHO decentralizes end user functionality and supports interoperability of distributed functions.

Although this Guide focuses on the needs of Client Partners, ECHO supports the following different, nonexclusive types of Partners:

- **Data Partners** Organizations that supply metadata representing their data holdings to the ECHO database

- **Client Partners** Organizations that participate by developing software applications to access the Earth science metadata in the ECHO database
- **Service Partners** Organizations that participate by advertising their Earth science-related services to the user community via ECHO, which maintains service descriptions in a Service Catalog (either special services, or services that are available as an option on a selected set of granules/collections) and support the user in ordering those services.
- **Extended Service Partners** Organizations that participate by providing a central location for registration, classification, and maintenance of Earth science services, interfaces, GUIs, and advertisements. ECHO addresses science user needs through a set of well-defined and open interfaces upon which the user community can build its own client applications. In this way, ECHO supports extendable, flexible user interfaces, allowing industry and the science community to drive the progress of available Earth science applications. For more complete information about client applications, refer to the companion piece to this Guide, the *ECHO 10.10*
- **Data Partners Guide** The ECHO approach allows users to build their own user interfaces to ECHO, rather than being limited to the data search and order system provided by NASA. For Data Partners, ECHO offloads the burden of providing the system resources required for searching and gives users the flexibility to support community-specific services and functionality. ECHO's interoperability features allow all participants to benefit from the distributed development of functions, again reducing dependence on NASA resources.

ECHO Benefits to Client Partners

To address the ECHO system vision, ECHO has responded to a set of system drivers, that is, reasons for upgrading. These drivers, derived from functional, organizational, and operational concerns expressed by the user community, determined the architectural approach and the types of technical solutions used in building the ECHO 10.10 system.

Ease of Participation

The primary goal of ECHO is to enable organizations to participate in making their resources and capabilities available to the Earth Science community. To facilitate participation by these organizations, ECHO has:

- Minimized the number of requirements that a partner must meet to participate
- Involved partners in the system's development cycle and requirements definition
- Selected metadata insert and update mechanisms based on current standard industry practice (for example, XML) that most databases can generate automatically
- Provided mapping capabilities to convert from one XML representation into another

Cost to Field

While aggressive in the capabilities it is targeted to support, ECHO minimizes the Cost to Field by continually evaluating performance and functionality against costs, for example, licensing of Commercial Off-the-Shelf

(COTS) applications, amount of custom code required, hardware platform requirements, and complexity of networking and installation.

Cost to Operate

Once fielded, ECHO seeks to minimize the cost to operate the system by making it easier to use, thereby minimizing the load on operations staff.

Extensibility

ECHO is being built with long-term extensibility foremost in mind. To enable emerging techniques and strategies for Earth Science research, ECHO has:

- Adopted a 'design for change' as a goal at the beginning of ECHO development
- Built in the capability to limit the impact of changes to the API on the configuration file, the service interface, and the business logic that implement the function
- Developed test tools to regression test large portions of functionality automatically overnight, so that when changes are made,

undesirable impacts can be discovered quickly

- Adopted a layered ECHO architecture to allow changes to one component of the system without affecting other components

ECHO Systems

There are three ECHO Systems that you, as a Client Partner, have access to:

ECHO Operations This is the current operational system for ECHO and is available to all users.

Location: <http://api.echo.nasa.gov/echo/index.html>

ECHO Partner Test This is an operational system used only by the ECHO partners where they can test their data and services prior to making the final changes in the operational system

Location: <http://api-test.echo.nasa.gov/echo/index.html>

ECHO Testbed This is a test system area used by partners and ECHO testers to test before changes to the ECHO system go operational

Location: <http://testbed.echo.nasa.gov/echo/index.html>